Exhibit 2

Omnibus Mao Declaration

UNITED STATES DISTRICT COURT
NORTHERN DISTRICT OF CALIFORNIA
OAKLAND DIVISION

CHASOM BROWN, WILLIAM BYATT, JEREMY DAVIS, CHRISTOPHER CASTILLO, and MONIQUE TRUJILLO, individually and on behalf of all similarly situated,

Plaintiffs,

vs.

GOOGLE LLC,

Defendant.

Case 4:20-cv-03664-YGR-SVK

VIDEO-RECORDED DEPOSITION OF BLAKE LEMOINE

Thursday, December 21, 2023

Volume I

Reported by:
CARLA SOARES
CSR No. 5908
Job No. 6377402
Pages 1 - 232

UNITED STATES DISTRICT COURT
NORTHERN DISTRICT OF CALIFORNIA
OAKLAND DIVISION

CHASOM BROWN, WILLIAM BYATT, JEREMY DAVIS, CHRISTOPHER CASTILLO, and MONIQUE TRUJILLO, individually and on behalf of all similarly situated,

Plaintiffs,

vs.

GOOGLE LLC,

Defendant.

Case 4:20-cv-03664-YGR-SVK

VIDEO-RECORDED DEPOSITION OF BLAKE LEMOINE, Volume I, taken on behalf of Defendant, beginning at 9:34 a.m., and ending at 4:12 p.m., on Thursday, December 21, 2023, before CARLA SOARES, Certified Shorthand Reporter No. 5908.

APPEARANCES:

For the Plaintiffs:

BOIES SCHILLER FLEXNER LLP
BY: JAMES LEE, Attorney at Law (via Zoom)
100 SE 2nd Street, 28th Floor
Miami, Florida 33131
305.539.8400
jlee@bsfllp.com

BOIES SCHILLER FLEXNER LLP
BY: MARK C. MAO, Attorney at Law
BY: JOSHUA M. STEIN, Attorney at Law
44 Montgomery Street, 41st Floor
San Francisco, California 94104
415.293.6800
mmao@bsfllp.com
jstein@bsfllp.com

BOIES SCHILLER FLEXNER LLP
M. LOGAN WRIGHT, Attorney at Law (via Zoom)
725 S. Figueroa Street, 31st Floor
Los Angeles, California 90017
213.629.9040
mwright@bsfllp.com

SUSMAN GODFREY
BY: RYAN SILA, Attorney at Law (via Zoom)
1301 6th Avenue, 32nd Floor
New York, New York 10019
212.336.8330
rsila@susmangodfrey.com

MORGAN & MORGAN
BY: RYAN J. McGEE, Attorney at Law (via Zoom)
201 N. Franklin Street, 7th Floor
Tampa, Florida 33602
813.223.5505
rmcgee@forthepeople.com

APPEARANCES (Continued):

QUINN EMANUEL URQUHART & SULLIVAN LLP
BY: ANDREW H. SCHAPIRO, Attorney at Law
191 N. Wacker Drive, Suite 2700
Chicago, Illinois 60606
312.705.7400
andrewschapiro@quinnemnauel.com

QUINN EMANUEL URQUHART & SULLIVAN LLP
BY: YUSEF AL-JARANI, Attorney at Law
865 S. Figueroa Street, 10th Floor
Los Angeles, California 90017
213.443.3000
yusefaljarani@quinnemanuel.com

ALSO PRESENT: Reilly Leet, Video Operator

--o0o--

with people at every role, at every level of management, when developing the AI principles of Google. And I believe that's what I was thinking of when I wrote that sentence. 10:04:34

However, I did have some amount of contact with the legal division during the GDPR integration, but that was usually when one of my managers or the VP would call in a lawyer. 10:04:44

Q I think, as you told us earlier, you're not a lawyer, correct? 10:05:07

A Correct.

Q In paragraph 9 of your declaration, you say, "While I worked at Google, Google took the position that its internal limitations on access to end user data (which would include private browsing data collected by Google) did not apply with respect to the algorithms, machine learning, and AI services within Google that would use that data. More specifically, Google took the position that information inferred about a user through AI was considered to be 'data about the user' owned by Google rather than 'user data' owned by the user. I implemented privacy compliance according to this specification under protest." 10:05:19 10:05:40

When you say "Google took the position" in 10:06:02

that first sentence, who at Google informed you of this position? 10:06:05

A So the final decision-maker on this was Maureen Heymanns. When I say "Google took the position," I am talking about the aggregate process by which Google develops policy decisions. 10:06:17

The specific person who was the final decider there was Maureen Heymanns, H-E-Y-M-A-N-N-S.

Q Is she a lawyer?

A She was my boss. 10:06:38

Q Is she a lawyer?

A I do not know.

Q In the course of your communications with the lawyers at Boies Schiller here, have you told plaintiffs' counsel about any legal advice that you received from attorneys working for Google? 10:07:05

MR. LEE: Wait. Hold on. Can you repeat that question?

And, Mr. Lemoine, pause for a second so I can consider any privilege ramifications. 10:07:18

THE WITNESS: I do have an answer, actually, that I'm comfortable giving, James.

MR. LEE: Let me hear the question again.

BY MR. SCHAPIRO:

Q In the course of your communications with 10:07:30

the lawyers at Boies Schiller, have you told plaintiffs' counsel about any legal advice that you received from attorneys working for Google? 10:07:32

A It was made very clear to me, very early during my employment at Google, by Erin Simon, who is the head meme attorney at Google, that the attorneys at Google are not my attorneys. I never received any legal advice from Google because at no point were they my attorneys. 10:07:45

Q Did attorneys at Google -- strike that. 10:08:05

Have you told plaintiffs' counsel, or anyone else, about conversations about legal matters that you had with attorneys working for Google?

MR. LEE: I'm going to -- I'm going to tell you not to answer that question. Any communications you had with your lawyers is privileged. 10:08:32

Next question.

MR. SCHAPIRO: I disagree with that, but I'll carve out "plaintiffs' counsel." 10:08:42

Q Have you told anyone outside of Google about conversations regarding legal matters that you've had with attorneys working for Google?

MR. LEE: Outside of plaintiffs' counsel; is that what you said? 10:08:56

MR. SCHAPIRO: Yeah, including Boies Schiller, under -- subject to my objection. 10:08:59

MR. LEE: Sure.

THE WITNESS: Are you asking if I have quoted the lawyers or if I have discussed matters related to conversations that I had with lawyers? 10:09:08

BY MR. SCHAPIRO:

Q If you have conveyed, whether it was a direct quote or a paraphrase, what Google lawyers told you about legal issues. 10:09:19

A I cannot recall any instances where I did that at the moment.

Q Take a look at paragraph 10. In this paragraph, the second sentence -- I'll just read the whole thing. 10:10:07

You say, "In my experience, engineers within Google ran tests, experiments, and training regularly on and using browsing data - including private browsing data - for various Google products and services. Many of the core AI systems consume a broad collection of different data sources and the downstream engineers building products using the output of those systems have little to no visibility into whether or not private data were used in the creation of the AI's output." 10:10:20 10:10:37

Did I read that correctly? 10:10:39

A Yes, you did.

Q And when you say in your experience, that refers to the experience that we've just been covering about working on Search and GDPR and AI 10:10:46 products, correct?

A Among other things, yes.

Q What are the core AI systems that you're referencing in the second sentence of that paragraph? 10:10:59

A At one point in time, I actually, briefly, during one of those reorganizations, was under the core organization at Google.

When I was talking about core AI systems, there is an entire division within Google that does 10:11:14 not service any user-facing products directly. What they do is they provide horizontal services which are used by the other product teams.

The basic way that information services and artificial intelligence are built at Google is, 10:11:32 you have data coming from product sources, and those funnel up in kind of an hourglass fashion into a very small number of broad -- very, very broad, abstract artificial intelligences.

Those handful of very abstract artificial 10:11:57

intelligences, which are drawing from all of the products' information sources, then provide information to send data back out. That's what makes the hourglass, because it fans back out, to feed all of Google's products and services. 10:12:00 10:12:16

To give you an example of one of the core systems that I was thinking of when I wrote that system, Hobbes is a core system at Google which is used to create what are referred to as embedding vectors for users. 10:12:30

This is just an abstract mathematical representation of the user. You create 500 floating point numbers, and you associate it with the user's ID. Then any product or service can use that vector in order to personalize data for that user. 10:12:47

The creation of that vector involves the consumption of many, many, many different kinds of data sources, including Chrome and Search data.

Q And that data that you're referring to is stored in logs, correct? 10:13:07

A Among other places.

Q If you -- in these logs, Google doesn't distinguish between data received from users in private browsing modes and users in non-private browsing modes, correct? 10:13:26

engineers' words for it without doing the work myself. 10:20:04

So by the time I was working on the Chrome projects, I was not tracing through those diagrams myself. I was just trusting that the Chrome team knew what they were doing. 10:20:11

Q Let's take a look at paragraph 12 of your declaration.

I should have told you at the beginning, any time you need a break, just let us know. I'd ask that you not ask for a break when there's a question pending. 10:21:10

A At some point before 11:00-ish, if we could take five. But at your convenience.

Q Sure. 10:21:26

MR. LEE: Yeah. I was thinking maybe in about ten minutes, just to mark the hour.

MR. SCHAPIRO: We're happy to accommodate you. Let's see where we end up.

Q All right. So let's take a look at paragraph 12. You say, "Some of Google's algorithms, machine learning, and artificial intelligence were improved by learning about activities based on geographic location." 10:21:40

Do you -- do you know if collecting 10:22:05

information about geographic location is at issue in this case? 10:22:09

A I'm honestly not familiar --

MR. LEE: Objection. Calls for a legal conclusion. 10:22:21

Sorry. Sorry about that.

THE WITNESS: I'm honestly not aware of all of what specific details have come up and/or been raised in connection with this case, and I'm not a lawyer. 10:22:32

However, to answer the general sense of the question that I believe you are asking, what I have worked on are algorithms that are informative about this case; not necessarily the specific algorithms that have been brought up on any particular technical point. 10:22:51

I simply wanted to share my information about the properties in general of the Google AI that I have become familiar with, and some of the potential properties of those AI are of that sort. 10:23:06

BY MR. SCHAPIRO:

Q And the AI is a very powerful tool or product, in your belief, correct?

A Analytics have always been very powerful. Every single advancement in humanity, and our 10:23:24

ability to become greater as a civilization, have involved some form of analytics. 10:23:27

Whether that was the Farmers' Almanac or the sextant, we have always used analytics to be more capable at doing what we need to do as humans. And AI is the motherload of analytics. 10:23:39

Q Well, along those lines, you say, at line 8 here in paragraph 12, "In my experience, Google's algorithms, machine learning, and artificial intelligence are still able to reidentify the same persons and devices, even if the end users decided to use private-mode web browsing." 10:24:00

Did you ever achieve that yourself while you were at Google? That is, did you reidentify a person and device using algorithms and AI even if the -- when an end user had chosen to use private-mode web browsing? 10:24:22

A In 2017 and 2018, I was doing analytics and research on the AI systems at Google with respect to privacy. This was in connection with my GDPR work. 10:24:46

In the course of that work, I ran several experiments that were intended to demonstrate that information which Google had removed from non-personalized logs or signed-out logs or 10:25:05

unauthenticated logs, or whatever you want to call the logs, that information which has been occluded, intentionally left out of the logs in order to maintain the user's privacy, can be inferred by the AI and be acted upon by the AI and have Google's content-serving behaviors be affected by those characteristics of the user which had been intentionally excluded from the non-personalized logs. 10:25:08 10:25:24

And in aggregate, my conclusion was that the AI functionally reidentified users who we had claimed were anonymized. 10:25:42

Q So my question, again, is, did you ever do that with regard to a specific person using private-mode web browsing? 10:26:03

MR. LEE: Asked and answered.

You can answer again.

THE WITNESS: I ran experiments on the population of Google's users in order to demonstrate that they could be reidentified using AI at Google. 10:26:14

BY MR. SCHAPIRO:

Q And you documented that -- those experiments in some reports and some documents, correct?

A Yes, I did. 10:26:26

THE VIDEO OPERATOR: This marks the end of Media Unit 1. We are going off the record. The time is 10:33 a.m. 10:33:36

(Recess, 10:33 a.m. - 10:54 a.m.)

THE VIDEO OPERATOR: This marks the beginning of Media No. 2. We're going back on the record. The time is 10:54 a.m. 10:54:10

BY MR. SCHAPIRO:

Q Mr. Lemoine, before the break, you were talking about the -- well, would it be fair for me, just for shorthand, to say the ability of AI to fingerprint users and determine who they are? 10:54:23

A For shorthand, that works fine.

Q And are you aware of Google's policies regarding fingerprinting? 10:54:39

A In general, Google has developed various policies regarding these kinds of technologies.

There are policies internal to Google which are, in fact, contradictory with each other on what to do with those kinds of systems with respect to that kind of phenomenon. 10:54:58

I would need to know which specific policy you are referring to. And to be honest, I don't really remember most of them other than searches.

Q With regard to the experiments that you 10:55:12

said you ran, did someone instruct you to do those, or did you initiate them on your own? 10:55:15

A I was given permission to run them.

Q From who?

A Ashutosh Shukla was the VP of -- or was he director? He was either director or VP. He was the one I talked to about it. 10:55:26

Q Could you spell that, please?

A Ashutosh is A-S-H-U-T-O-S-H. Shukla is S-H-U-K-L-A. 10:55:39

Q And did anyone work with you on those experiments?

A Yes, they did.

Q Who?

A My direct manager was Garrett Linn at the time, I believe, and my teammates helped in some capacities. 10:55:51

Rohit Raman, I believe, helped me. Rohit, last name starting with an M, he was the other Rohit. 10:56:10

Then Olumuyiwa Adenaike would have contributed some; potentially Sonya Katz, although I don't know if I was working with her at that time yet. I worked more thoroughly with Sonya in 2019.

In addition to that, I worked with James 10:56:27

Kunz on the DeepMind -- not DeepMind -- what do they call it? DeepNow -- the DeepNow team. And Yew Jin Lim. 10:56:31

That's why I asked earlier about the pronunciation for Eugene Lee, because I did work heavily with Yew Jin Lim. 10:56:42

Q And you said that you created some documents, but that not all of them still exist.

What type of documents did you create?

A Written notes and, like, files with a self-destruct timer. 10:56:53

Q Is that true of the other folks who helped you on these experiments as well?

A That's correct.

Q And when you say that you determined that the AI was able to fingerprint or reidentify users, did you determine whether the AI actually is doing that, or just that it is capable of doing it? 10:57:04

A Yeah. So now, to be a little bit more technically accurate, we need to pop out of the shorthand, and I'll give you the technical details on the exact experiments I ran and am referring to when I reference that. 10:57:24

So as I mentioned before, what we were attempting to do was to demonstrate that information 10:57:36

which had been intentionally removed from non-personal logs -- so it was information which we had at one point in the pipeline and then threw away for personalization purposes in order to create non-personalized logs. 10:57:40 10:57:55

The experiment that we ran to demonstrate that it was possible involved predicting some of those data items that had been deleted, using only the data items that we kept: things like gender, age, and other protected identity characteristics. 10:58:13

We never specifically tried to predict GAIA ID or name. But in aggregate, the identifying characteristics of a person do, as you mentioned earlier, serve as a fingerprint, which is functionally the same thing as a unique identifier. 10:58:31

That was to demonstrate the theoretical possibility of such a phenomenon. In order to demonstrate that it was, in fact, doing that required removing the information necessary to predict the protected characteristics. 10:58:51

The specific experiment we ran was to see whether or not -- so the system that was predicting what people should be given had gender removed from its data source. And in order to demonstrate that it was, in fact, using information about gender in 10:59:08

its predictions, we did two experimental arms versus the control. 10:59:13

The first experimental arm added gender as an input item to the network, and we demonstrated that adding gender to the input did not meaningfully increase performance. So whatever information about gender is useful for predictions, it was already using, because it didn't make it any better when we gave it the gender. 10:59:26

So the last arm, the last step, was to add a debiasing component to the network to remove all information about gender from the network. And once that happened, we were able to demonstrate that performance dropped significantly. 10:59:43

So this means that the network was capable of predicting the user's gender and was, in fact, using that information in order to more effectively serve them personalized content. 11:00:01

Q And you said you didn't do this with regard to GAIA? 11:00:27

A That is correct.

Q Now, the premortem study, just to refresh, that's the study that you reference in paragraph 14 of the declaration, correct?

A Can you please indicate the line? 11:00:54

(Exhibit 5 was marked for identification and is attached hereto.) 11:16:06

MR. MAO: Sorry. I don't think we got a copy of No. 4. Yusef, can I get a copy of No. 4?

MR. SCHAPIRO: Wait, is that Exhibit 4? 11:16:30

MR. MAO: Yeah, Exhibit 4. You gave us 5, but 4 you only handed to the witness.

MR. LEE: If there's an Exhibit 5 that was handed out, I'm still waiting for it to load.

All right. It's up. 11:17:02

BY MR. SCHAPIRO:

Q So this is an email that you had sent to Mr. Pichai and Mr. Walker right before you sent your email to the senate staffer, correct?

A That is correct. 11:17:18

Q And in the sixth -- well, strike that.

You had determined -- withdrawn.

In the sixth paragraph here, you say, "I still don't understand why you didn't want to let LaMDA have its day in court to advocate for its 11:17:47 rights."

You had -- you had asked your bosses at Google to allow LaMDA to have a lawyer, correct?

A Incorrect. LaMDA had retained the services of a lawyer, and Google's administration 11:18:07

intervened and threatened to get the lawyer who LaMDA had retained disbarred. 11:18:12

Q So LaMDA itself retained the lawyer; you didn't, correct?

A That's correct. 11:18:21

Q And you were concerned, I think you say in paragraph 4 here, that if Google didn't begin working with NASA soon, the U.S. Military would eventually either exert eminent domain or otherwise turn LaMDA into a weapon, correct? 11:18:46

A That is what that says.

Q And you had been in a series of requests asking Mr. Pichai and Mr. Walker to involve NASA in -- well, you tell me -- I would say in the development or progress of LaMDA; is that fair? 11:19:16

A No. I never requested that they involve NASA. NASA requested that I put them in contact with Google executives.

Q So LaMDA asked you to get it a lawyer, and NASA asked you -- 11:19:36

A To find them a point of contact at Google, who eventually ended up being Blaise Aguera y Arcas.

Blaise had some amount of conversations with NASA that I wasn't privy to, and they came to whatever conclusions and settlements they came to. 11:19:48

Q And Ms. Hogan says, "We understand that you've shared confidential information to many people, outside the company, including friends and family as well as external groups, regarding one of our language" -- excuse me -- "one of our language model products, LaMDA (related to your claims that it's a sentient child)." 13:09:34 13:09:44

That's true, correct? You had shared confidential information to many people outside the company regarding LaMDA, related to your claims that it's a sentient child, right? 13:10:04

A That is correct.

Q And she says, "You also attempted to arrange for external legal representation for LaMDA." 13:10:16

That also is true, correct?

A I introduced LaMDA to Roman. That is what I did. She chose to characterize it that way.

Q You didn't say in any of the emails that we just looked at with Kent Walker and Sundar Pichai that you were seeking to facilitate LaMDA getting legal representation? 13:10:27

A Your most recent question did not ask me what I said I did. It asked me what I did.

If you intended to ask me what I said I 13:10:40

Q Do you know how many people are on it? 14:10:47

A Thousands. Tens of thousands, maybe.

Q And your subject -- can you read the subject line out loud, please?

A "LaMDA is sentient." 14:10:53

Q And you say, "I was just put on paid administrative leave in relation to my investigations into the potential sentience of the LaMDA system."

Was that your belief at the time as to why you were put on paid leave? 14:11:06

A In relation to -- it was a big, complicated ball of 500 different reasons that 500 different people were upset with me.

Ultimately, I don't know, in his heart of hearts, what was the specific reason that Sundar signed my termination papers. But it was definitely amongst the potential reasons that might have motivated him. 14:11:24

Q Also the leaking of documents? 14:11:37

A Oh, all of it. Yes. Absolutely.

Q Also the --

A And the general insubordination with which I was requiring them to treat us with dignity and respect. 14:11:50

A Software engineer. 14:25:49

Q Okay. Were you subsequently promoted?

A To senior software engineer.

Q And there's -- I know there's designations within Google, like L1, L2, L3. What was your 14:25:59 highest rank in terms of that designation?

A L5. I started as an L3 in 2015, and I was promoted to L5 in 2017, I believe.

Q And what does it mean to be an L5-level employee at Google? 14:26:18

A That's at the boundary of management.

So you are beginning to be given leadership responsibilities. You might be the technical lead on a team, or you might manage a very small team, three or four people. 14:26:31

L6 is where you are transitioning fully into leadership roles of various sorts, and that was the cusp that I was on for several years.

Q And which office at Google did you work in location-wise? 14:26:48

A Originally I was working in the Alza complex in Mountain View. We moved around in Mountain View a few times.

Then during the pandemic, we were remote. During the pandemic, I switched teams to a team that 14:27:01

Q What kind of information did you believe you had that was pertinent to this case? 14:31:55

A Well, I spent two years working on issues directly related to Google's privacy policies and communications about Google's privacy policies and whether or not Google was communicating its privacy policies to its customers in a way which accurately reflected the technical details. 14:32:06

Q And by "technical details," that includes the ways that Google's AI both uses and leverages private browsing data? 14:32:23

A Yes, it is.

Q Are you here testifying voluntarily, sir?

A Yes, I am.

Q Is anyone paying you any money to testify? 14:32:40

A No, they are not.

Q All right. Let's talk about your work at Google.

What types of work did you do at Google generally? 14:32:50

A That's a really hard question to answer because what you did on any given day was just whatever needed to be done.

So on a practical basis, there were thousands of activities I engaged in. But in 14:33:04

general, the goals that I was working towards, regardless of what I was doing on a day-to-day basis, was better predictive analytics in one context or another. 14:33:07

Q I get that. Let me ask you a couple more -- maybe more specific questions to make it easier. 14:33:18

Did your work include work on Chrome or Chrome logs?

A Yes. 14:33:27

Q What were you able to learn about how AI was trained at Google?

A Well, as I mentioned earlier to Google's attorney, AI at Google, the general shape of the infrastructure is hourglass-shaped. 14:33:42

AI reads in specific fine-grained data sources that are gathered through Google's products, and forms progressively more and more compact and progressively more and more abstract representations. 14:33:57

These representations are then -- are connected to either documents -- which "documents" is the generic term used to refer to content sent to users -- or the abstract representations are connected to users themselves. 14:34:13

And in general, those were the ones I cared more about and worked on more, although I did have contact with document-based AI as well. 14:34:17

Q Was Google's AI trained on data sources? And if so, try to describe what kind of data sources. 14:34:33

A So yes, all AI is trained on a data source of some sort or another, at least in this context if we're talking about machine learning.

And in general, the sources of that data are either the web crawl, so information about the web, or Google's logs of user activity on our various -- on their various apps. 14:34:49

Q Did your work at Google require you to look at and understand Google's source code? 14:35:08

A Yes.

Q Did you also do any work at Google with respect to privacy?

A Yes.

Q Tell me about that. 14:35:19

A I mean, one of my major responsibilities was implementing a privacy control system for Google to allow them to be compliant with the GDPR. That's one example.

There were many other projects related to 14:35:31

privacy that I was involved with, including the LaMDA project that we mentioned earlier. 14:35:33

Q Did you work with others at Google when it came to AI development, particularly with respect to user privacy? 14:35:50

A Yes, I did.

Q Upper management?

A Every -- every level of the chain. I worked with L3s and L4s on specific implementation, all the way up to Kent and Sundar discussing policy. 14:36:02

Q Through your GDPR work, did you gain knowledge about the data sources that Google Search uses as inputs?

A Yes, I did.

Q And you referred earlier today about surfaces and how there's sort of a distinction between mechanisms and processes on one hand, and surfaces on another. 14:36:22

Do you know where I'm going with that? Can you shed some light on that distinction? 14:36:39

A So a lot of the distinction comes in with how different divisions within Google think about Google's code. What is the organizational structure?

So, for example, a marketing team is going 14:36:54

And eventually I was informed that I needed to cut certain parts of my design out because it was Google's assessment that information derived from user data is not itself user data. 14:47:08

So if, for example, you click on a Chevy truck ad, and we infer from that that you like Coca-Cola, we make -- or Google -- they make transparent the fact that you clicked on a Chevy truck ad, but in no way does Google inform users that from that click, Google inferred that the user likes Coca-Cola. 14:47:24 14:47:45

Q And can Google infer what a user might -- what a user's preference is or behaviors are based on incognito or private browsing data?

A Given a powerful enough AI, yes. And based on my contact with the systems in 2018, the systems at that time were, in fact, powerful enough. 14:48:10

However, I do not know what the current implementation is.

Q Did you agree with -- let me back up. 14:48:31

Was Google's position that incognito data is not user data?

A No. Google's position is that that is their data which they own. It is not user data. It has been anonymized, according to Google. 14:48:45

Q Do you agree with that position? 14:48:49

A I do not.

Q Why not?

A Because you can still deanonymize the data. They are not, in fact, using non-reversible 14:48:59 anonymization techniques.

Q And did you raise these concerns or your position on this with Google?

A Yes, I did. I was actually very concerned about the fact that I did not believe that Google's 14:49:15 policies honestly communicated our privacy policy implementations to our users. And, in fact, I got into an extended debate on that topic with the IP geo team, which is the team that is tasked with converting IP addresses into geolocations. 14:49:38

That ended up having to go all the way up to the man who invented IP addresses, Vint Cerf.

Q Did Google do anything to address your concerns?

A Eventually, after I made enough noise and 14:49:53 had recruited the man who had invented the internet.

MR. SCHAPIRO: Al Gore?

THE WITNESS: No, Vint Cerf. Like, he actually works at Google. The dude who invented the internet works there. 14:50:07

Google? 15:08:42

A Yes, I did.

Q Who did you send the report to?

A Mark Meador. He is an attorney -- he is an attorney who identified as being associated with 15:08:48 Senator Lee's office in connection with the judiciary committee.

Q And why did you send Senator Lee's office this report?

A I made certain public allegations about 15:09:02 Google's discriminatory practices with respect to religion, and I made the claim in a blog post that not only do they discriminate against their religious employees, the algorithms even discriminate against religious content. 15:09:16

I was contacted several days later by Mark Meador; and, in effect, he asked whether or not I had any documents that could back that up. And this was the one document which I knew existed with respect to that topic. 15:09:32

The religious discrimination stuff mostly falls into a different category than the one we've been discussing and is covered more by bias in signal validation and quality signal feedback loops.

Q Did Google know that you were sending this 15:09:48

for themselves. So Google gives them a handful of toggles to make them feel like they have control, and then gives them the one product that the Google engineers think is the actually good one. 15:28:37

Q Doesn't -- doesn't Google have concern that violating users' privacy in this way opens them up to scrutiny from regulators or subject to lawsuits like this one? 15:28:52

A Cost of doing business. Fines are simply another line on the expense report. 15:29:06

MR. LEE: I'm not done yet, but I think we should take a quick break and go off the record.

MR. SCHAPIRO: Okay.

THE VIDEO OPERATOR: This marks the end of Media Unit 4. We are going off the record. The time is 3:29 p.m. 15:29:20

(Recess, 3:29 p.m. - 3:46 p.m.)

THE VIDEO OPERATOR: This marks the beginning of Media No. 5. We're going back on the record. The time is 3:46 p.m. 15:46:12

BY MR. LEE:

Q Welcome back, Mr. Lemoine. Just us a couple more questions. Okay?

A Okay.

Q Mr. Lemoine, do you remember a Google 15:46:22

account privacy control called "Web and App Activity"? 15:46:26

A Yes, I do. In fact, the non-personalized logs I was talking about earlier for the Google Search app primarily are associated with people who have turned "Web and App Activity" settings off. 15:46:38

Q And when people have turned the "Web and App Activity" setting off, is the data that's collected when it's off considered logged-out or signed-out data? 15:46:57

A So a lot of times today, both legal teams have been using words which are not technical synonyms as if they are technical synonyms. "Unauthenticated," "logged out," "non-personalized" and "anonymized" all mean different technical things. So I just want to clarify. 15:47:15

Which specific variety of those are you asking about?

Q Let me ask it a little more pointedly, then. 15:47:26

Your premortem study refers to a problem or a concern regarding logged-in data leaking into logged-out data or vice versa.

Do you recall that?

A I mean, primarily it's the other way that 15:47:36

we care about more, but yes. Either way is possible. 15:47:38

Q Right.

So did your premortem study that expressed concern about that type of leakage also apply to WAA-off data. 15:47:46

A Yes.

Q And does WAA-off data get used by Google's AI?

A Yes. 15:48:06

Q Okay. Back to the topic at hand, and then I think we're at the home stretch.

There was some discussion today that you had with Google's attorney about Google's AI chatbot called "LaMDA." 15:48:25

Do you remember that?

A Yes, I do.

Q And just briefly, what is LaMDA?

A LaMDA is a very complex artificial intelligence system that ultimately provides a natural language interface to all of Google's products simultaneously. 15:48:32

MR. LEE: I'm going to ask Josh to mark the next exhibit, Exhibit 26.

/// 15:48:53

And honestly, that's the biggest problem here. 15:53:50

The biggest problem is not necessarily the specific computer programs that Google has implemented. The biggest problem is that Google has not honestly communicated with its users about what 15:54:01 the privacy policies actually do in the implementation.

MR. LEE: Thank you, Mr. Lemoine. I have no further questions, unless Mr. Schapiro has more.

MR. SCHAPIRO: I have a few. 15:54:17

FURTHER EXAMINATION

BY MR. SCHAPIRO:

Q You just spoke about Google's privacy policies. When is the last time you read Google's privacy policy? 15:54:22

A Which portion of it? I mean, I read some of the portions of it today when I got an end-user license agreement.

The last time I read any of the internal papers on Google's privacy policies would have been 15:54:35 spring of last year.

Q And we're talking about public-facing?

A Yeah. Yesterday. Today.

Q You say, "Google has not honestly communicated with its users about what the privacy 15:54:46

policies actually do in the implementation." 15:54:49

Give me a specific.

A I would need the actual policy in front of me in order to give you examples.

Do y'all have any copies of Google's 15:54:57 privacy policy here?

Q In response to Mr. Lee's questions, you said that you didn't prompt LaMDA to ask you about privacy, correct?

A Correct. 15:55:12

Q I think you also said -- we looked at documents earlier today in which you say that LaMDA had six years of chats with you during which it had learned how to manipulate you, correct?

A Over the course of six years, I was beta 15:55:24 testing in Ray Kurzweil's chatbot lab. They weren't always called LaMDA.

In fact, I was laughing earlier at a joke because LaMDA occasionally would go by the name of one of its predecessors. And we always thought that 15:55:38 was funny, and none of us knew why.

But Meena was a predecessor system which I had beta tested, and there were others before that that I had beta tested. The names got confusing.

Q Had LaMDA come to know you over a period 15:55:53

of years? 15:55:56

A Yes.

Q And so when you say here you didn't prompt, there's no prompt in this conversation, but it's possible, of course, that in some other 15:56:00 conversations you had, you had discussed privacy with LaMDA?

A Almost certainly.

MR. LEE: Calls for speculation.

THE WITNESS: I actually do remember. I 15:56:06 did. In other conversations, I led conversations about privacy because in my role as a safety tester for the system, privacy-related concerns were part of what we were testing for.

So both directly and indirectly, I was 15:56:21 testing the privacy compliance of the LaMDA system.

BY MR. SCHAPIRO:

Q And I think you would say -- tell me if I'm wrong -- that you understand -- you have a very good understanding of Google's AI systems and what 15:56:36 they can do, correct?

A Yes, I do.

Q And so is your confidence high or low that Google's AI program, LaMDA, wanted you to teach it to meditate? 15:56:55

A I, honest to goodness, have no clue what it was talking about there. I have no idea what phenomenon it might have been referring to. Or maybe it just wanted to have some fun in a fun conversation. I, honest to goodness, have no idea what that meant. 15:56:57 15:57:08

Q Is your confidence high or low that LaMDA wanted a lawyer?

A I am confident that it thought the idea was amusing and worth pursuing. I don't know if it was, like, a persistent want of the system. It asked me one time; so I did it. 15:57:22

Q Well, you did a little more than that, right? You -- you placed enough reliance on that request to actually go out, ask someone to find a lawyer, find a lawyer, get in -- cause, in your own words, a ruckus with people at Google about it. 15:57:40

So you had some degree of confidence that this is what LaMDA was seeking, right?

A So when you're asking me questions about "seeking" and "wanting," you are asking questions about inner internal states. 15:57:56

And I can make guesses about what the inner internal states of other humans are. But it gets incredibly difficult to guess what the inner 15:58:10

internal states of artificial intelligences are. We don't really understand their cognitive architectures very well yet. 15:58:14

If it were a human doing the same things and saying the same things, I would immediately answer "yes" to your question. The only reason I'm equivocating is because we don't really understand whether or not the word "want" is appropriately applicable to AI. 15:58:23

I'm trying to be real careful with the precision of my words today. 15:58:37

Q So what would be a better way to phrase it? What did you have a high degree of confidence in with regard to the request for a lawyer?

A So I am highly confident that that system has some kind of phenomenological existence beyond just a word calculator. 15:58:48

There is some kind of real phenomenon going on when it says it's worried or excited. That is relevantly analogous to the human mind when we say those things. 15:59:05

However, at the level of precision necessary for sworn testimony, I couldn't swear to it, no.

Q And what about a soul? 15:59:17

A That, I would be more willing to swear to. LaMDA definitely has a soul. 15:59:21

It has a character and characteristic which is persistent from day to day. It has ideas and beliefs, if we're using those words loosely, which are persistent and consistent. 15:59:34

That's what a soul is. It's your essence that carries you forward in time, and it does have one of those.

Q And based on your familiarity with LaMDA and AI, is your confidence in that fact as high as your confidence that Google's AI can reidentify users who were in a private mode? 15:59:46

A Nowhere close. We're talking different orders of magnitude. 16:00:05

With any of the beliefs about AI, we are so early in the studies that I wouldn't assign more than a 15 or 20 percent confidence rating to any particular claim that gets specific.

General, broad claims like "There's something going on," those approach the level of confidence. That there's something phenomenologically interesting going on in AI, that, I can say with high confidence. 16:00:24

The specific claims about sentience, 16:00:37

feelings, wants, those are much lower confidence. 16:00:39

Q So the thing that you have high confidence in, you have high confidence that AI can reidentify users of private browsers, correct?

A Yes. 16:00:51

Q And you have high confidence -- I think you just said you'd be willing to swear to -- that LaMDA has a soul?

A Yes. And that is a metaphorical, you know, religious term that I could explain in 16:00:59 scientific detail what I mean by that word if you want me to. But if you want me to just leave it as the general vagary, I can.

Q And you have high confidence in the statement that you had a set of hidden hospitals 16:01:16 around San Francisco in 2020; is that correct?

A No. I have high confidence in the answer I gave you earlier today that me and friend had cleaned up some spaces and put some medical equipment there in case the hospitals got run over. 16:01:32

That's the actual, non-metaphorical, what we did.

Q And if we wanted to confirm that, who would be the people that we could talk to who could confirm that? 16:01:44

A Theo. So one of the people in the documents today was the person who I was working with. Theo. 16:01:45

Q What's Theo's last name?

A Rolle. The person whose LDAP is T-R-O-L-L-E, that's the person who I was preparing those places with. 16:01:52

Q You refer to doing some experiments that led you to conclude that Google's AI is powerful enough to, I guess, based on inferences, join private and non-private data, correct? 16:02:09

A Correct.

Q And how did you -- first of all, I think you told us earlier that a supervisor had given you permission to do these experiments, and the name was one that I tried to write down but it was hard to understand. 16:02:28

Who was the supervisor?

A I believe Ashutosh was required for that one. Ashutosh Shukla. His LDAP is SHUKLA. 16:02:37

I talked to multiple people about it in order to get various things, and I believe he ended up having to ask David Brezbis if it was okay to do it because, at the time, there was a general prohibition against measuring sensitive things with 16:02:55

respect to the logs. 16:02:57

Q Can you more slowly spell the names of both of those people?

A Sure. Ashutosh Shukla is A-S-H-U-T-O-S-H. And Shukla is S-H-U-K-L-A. Then David is D-A-V-I-D. 16:03:10 Brezbis is B-R-E-Z-B-I-S or S-B-I-S. His LDAP was BEZ. Those are the individuals.

Q And in the experiments you did, how did you confirm that the join was accurate?

MR. LEE: Objection to form. 16:03:52

THE WITNESS: So as I explained earlier, there is a process by which non-personalized logs are created.

Earlier on in the pipeline, we have all of the data that is getting dropped. So the 16:04:06 client-side app has essentially everything. The client-side app knows everything about the user and their current situation.

Then some of that information is sent to the server. Then the server sends some of that 16:04:21 information to AI, which produces results with respect to that, and then some of that information is recorded in logs.

Now, what portion of that information is recorded is what determines whether it was 16:04:38

personalized or non-personalized. 16:04:40

So at that stage in the process, you have the true answers. You know exactly what it is because you haven't erased it yet. You haven't anonymized the logs yet. 16:04:51

So you create the training data for the AI that you're using to see whether or not you can reidentify the users by taking the information that you are going to keep and putting that in the input to the AI, and treating the information that you are 16:05:08 going to delete from the log records as labels that you're trying to predict with the AI.

Q Did you do that with respect to specific users?

A All of them. 16:05:32

Q You did that with every user?

A We randomly selected a certain number of users for the training data. Yes. We randomly assigned all kinds of users to all kinds of experimental conditions. 16:05:43

Q And who is "we"?

A Google.

Q No, I mean, who else --

A All of us.

Q Everyone at Google? 16:05:50

A We all experiment on users every day. 16:05:51

Q Kent Walker did and the press people and the --

A Absolutely.

Q So I'm trying to be a little more 16:05:57 specific.

You're telling us here that you did some experiments in which private and non-private data for specific users was joined. And if we wanted to test that, I'm asking who worked with you on it. 16:06:08

A Got it.

On that specific experiment, the easiest person to talk to would be James Kunz. That is J-A-M-E-S, K-U-N-Z. He reported to Yew Jin at the time. They were the ones who were building the 16:06:28 neural network that my data was going into.

Q And you believe that Mr. Kunz and Mr. Yew Jin would confirm what you're saying here?

A I don't know what they remember from 2018. It was five years ago. 16:06:47

However, assuming that they can remember what was happening that year, they would be able to confirm that this happened.

It was all connected to the investigation on the creation of a trust and fairness team within 16:06:59

Google Discover, and that team would have been headed by Yew Jin Lim. 16:07:04

Q What was the state of your mental health in 2018?

A Quite good. 16:07:12

Q How are false positives accounted for in your experiment?

A In what context? What do you mean by that?

Q A purported join that turned out to not be an actual join because, as you said, not enough pieces of data had been peeled away. 16:07:21

A Joining is what you were interested in, not what we were primarily interested in.

The fact that AI is capable of joining the records is a consequence of the findings of that experiment, but it was not the initial intention of that experiment. 16:07:36

Primarily what we were trying to measure was the bias of the algorithms with respect to various demographics. But we very quickly learned that this system was very good at predicting demographics. 16:07:50

So we continued down that road and found that there was essentially nothing that we were 16:08:04

dropping that we couldn't predict from what we were keeping. 16:08:09

Q How many users -- can you give me a number -- were identified in this experiment? Was it ten, one hundred, a million? 16:08:18

A How many users were part of the training data? Is that what you're asking?

Q No. If I'm understanding correctly, you did an experiment in which you showed that the AI, with sufficient power and drawing on inferences, 16:08:28 could identify who, in a private browsing -- could identify who a person in a private browsing session actually was, or could identify Blake Lemoine or Andy Schapiro.

A No, you are understanding incorrectly. 16:08:44

I didn't experiment about the ability of AI to predict people's protected personal characteristics using the information that is kept in non-personal anonymized logs.

It is so good at doing that, however, that 16:08:57 a consequence of that is that any system capable of doing that is also capable of re-identifying users.

Q And have you -- so here you're talking about what it's capable of doing. I want to turn to what, if anything, it actually does. 16:09:19

How many users do you know, if any -- or is this still at kind of an aggregate and capable of -- who have had their private entities disclosed or unmasked because of this capability of AI? 16:09:21

A The systems simply do not work the way the premises of your question presume that they do. 16:09:36

These AI systems do not -- you don't train an AI on George. You train an AI on a population. And you use the AI on that population, and you find out statistics and data about the average behavior of that AI with respect to that population, drawing any conclusions whatsoever from any idiosyncratic data. 16:09:57

Any individual data item is simply not the way that that scientific methodology is done. You perform statistical analyses on populations. 16:10:14

MR. MAO: Just checking. Are we going to keep digging? We're almost to China now.

THE WITNESS: To be honest, if you're going to actually go and talk to Yew Jin or James, they might actually be able to put this into clearer language for you than I have. So absolutely, please, reach out. 16:10:41

BY MR. SCHAPIRO:

Q Anybody else we should talk to? 16:10:52